



PACKET PRIORITIZATION PROTOCOL FOR A LARGE-SCALE, HIGH SPEED COMPUTER NETWORK

[illegible]

1 **PACKET PRIORITIZATION PROTOCOL FOR A**
2 **LARGE-SCALE, HIGH SPEED COMPUTER NETWORK**

3
4
5 **BACKGROUND OF THE INVENTION**

6 **1. Related Applications**

7 This application is a Continuation-In-Part of and claims priority to our co-pending
8 patent application Serial No. 09/500,721, filed on February 9, 2000, entitled LARGE-SCALE
9 HIGH SPEED COMPUTER NETWORK AND METHOD OF IMPLEMENTATION AND
10 OPERATION, which is incorporated herein by reference.

11
12 **2. The Field of the Invention**

13 The present invention relates to computer communications networks. More
14 specifically, the present invention relates to methods of chronologically prioritizing time-
15 sensitive data transmissions from a switching station.

16
17 **3. The Relevant Technology**

18 Computer technology is breaking barriers to inter-personal communications at an
19 amazing rate. Already, it is possible to communicate almost instantaneously with anyone in
20 the world that has a computer and a telephone line. Computer networks, such as the Internet,
21 link individuals and various types of organizations in world-wide digital communication The
22 Internet has almost unlimited promise for communications advances, but is limited by an
23 overburdened and somewhat unsuited transmission medium.

24 In addition to the Internet, businesses, educational institutions, government agencies,
25 and other similarly related entities also communicate over much smaller-scale networks, such
26 as local area networks (LANs) and wide area networks (WANs). These small-scale

1 networks, particularly LANS, operate at much higher speeds than the Internet, but are
2 expensive to operate at large scales. Thus, a large gap exists, between the scope of coverage
3 and speed of operation of the global, but relatively slow, Internet and the faster but more
4 limited LANs and WANs. It would be advantageous to close this gap with larger-scale
5 networks that operate at speeds close to that of LANS.

6 Several barriers exist to filling the gap between current limited coverage networks
7 and the Internet. One such barrier is the "last mile" dilemma. That is, the Internet runs at
8 very high speeds over its backbone, but slows down considerably over its localized
9 connections. Generally, the Internet relies upon standard telecommunications industry lines
10 and switching equipment for this last mile. This infrastructure is designed for telephone
11 communications, and is not well adapted to the packetized communications of digital
12 networks. A dilemma lies, however, in replacing the telephone infrastructure with
13 transmission mediums more suited to digital communications. It is currently considered
14 prohibitively expensive to connect high speed communications lines down to the individual
15 users of the Internet.

16 This fact, together with the general congestion of the Internet in general leads to a
17 substantial slow down of Internet communications. It also limits the deployment of
18 intermediate types of networks. A further barrier to the implementation of networks of
19 varying scopes and to the new introduction of new paradigms for network communication
20 comes in the form of financing. Such developments using current technology would be
21 prohibitively expensive. Who is going to pay for this infrastructure?

22 Accordingly, a need exists for an intermediate sized network to close the gap between
23 the world-wide Internet and current relatively small scale networks. Preferably, such an
24 intermediate sized network operates at speeds similar to those of LANS, coverage both in
25 geographical area and diversify of user type. Additionally any solution to this problem
26 should also address financing of installation and should overcome the last mile dilemma.

1 New technologies for achieving such a new paradigm in computer networking are similarly
2 needed.

3 In addition to the lack of larger scale, high-speed networking, prior art networks of
4 every size have additional problems. Many of these problems result from the way in which
5 switching is carried out by known networks. Switches are simply junctions for multiple
6 communication lines. A "data transmission" is simply an analog or digital signal sent from
7 an origin to a destination. Bundled data transmissions, or "packets," arrive at an incoming
8 port of the switch, and are routed to the proper outgoing port to reach their destination.
9 (Although each port is capable of two-way communication, the port through which a packet
10 arrives is designated as the incoming port, while that port through which it will exit is the
11 outgoing port.) A data transmission from one computer to another may pass through several
12 switches, depending on the size of the network involved. Full-duplex, switched networks
13 are generally far faster than their half-duplex, unswitched counterparts.

14 However, a special problem arises when multiple packets simultaneously arrive at
15 a switch through different incoming ports, and all of the packets must go through the same
16 outgoing port. Since a line is only capable of conveying a single packet at a time, one packet
17 will be sent while the rest wait. Current networking systems possess significant drawbacks
18 in that they entirely fail to prioritize, or prioritize improperly, the order in which the packets
19 are transmitted.

20 This has many undesirable effects. Since the switches are typically utilized in a
21 branching network, many more switches may be downstream from the outgoing port. The
22 switch itself is unaware of what type of destination any packet is sent to. The destination
23 may be a server hosting many users simultaneously, or it may be a single home user.

24 As a result, people waiting for critical communications are forced to wait for other,
25 less important traffic. For example, a company may have a large number of employees
26 receiving e-mail through a server on the network. The e-mails may contain important

1 instructions, information, questions, etc. that should not be delayed. However, if the e-mail
2 is routed through the same switch as a large file download requested by a computer near the
3 mail server, i.e., at the same outgoing port, the e-mail traffic may be slowed down by waiting
4 for the file download. This occurs even though delays are inconsequential for the download,
5 which will require several minutes in any case. Similarly, a number of computer users
6 performing research over the Internet, using a variety of different sites, may be slowed down
7 by a single user transmitting real time game data to another user.

8 No previously known system provides a sufficient solution to this problem. Simply
9 sending packets through the port in sequential, cyclical form, or "round robin" form, provides
10 equal time to each communication through the switch, and causes the problems described
11 above. Giving priority to the heaviest user, i.e., the destination that has received the most
12 packets, is inadequate because the volume of data is not proportional to its importance.

13 Thus, there is a need, unfulfilled by the prior art, for a new method for prioritizing
14 transmission of packets from a switching station. The method should preferably prioritize
15 transmission according to the destination that is receiving the most important, i.e. time
16 critical, information, while avoiding entirely blocking other destinations for lengthy periods
17 of time. In addition, hardware and suitable data structures are needed for carrying out the
18 method described above.

19 Another problem with known networks is broadcasting. Broadcasting occurs when
20 a packet is sent to an unresolved destination. Communications over the Internet often take
21 place on the third, or network layer of the ISO/OSI model, which is the Internetwork, or IP
22 layer, of the TCP model. Transmissions may be addressed to a certain IP address, but the IP
23 address is a property of the network, and may not be the same for a given device every time.
24 Internet service providers (ISP's), for example, will often assign a temporary IP address to
25 each individual dialed up computer.

26

1 In order to successfully route a packet to the proper device, a switch must have access
2 to the hardware, or MAC address of the device, which is unique to each individual network/
3 interface card (NIC) that connects a computer to the network. The MAC address
4 corresponds with the second layer of the ISO/OSI and TCP models. A computer sending a
5 transmission may not always have access to the receiver's hardware address.

6 Thus, the sending computer sends an address resolution protocol (ARP) broadcast,
7 or packet without a specific MAC address destination, which will then be propagated to
8 multiple computers. The ARP broadcast contains a designated IP address for the destination
9 computer, and acts as a request for a requested MAC address of the computer that has that
10 IP address. The computer that has the IP address responds by sending a packet back to the
11 origin of the broadcast, with its MAC address included in the packet. The computers can
12 then communicate directly over the network without broadcasting to other users.

13 The problem with ARP broadcasting is that it creates a great deal of unnecessary
14 traffic on a system. The ARP broadcast itself typically does not contain a great deal of data,
15 but it must be transmitted to many computers, thereby occupying a great deal of bandwidth.
16 Even if a receiving computer's MAC address is resolved by one transmitting computer on
17 the network, another transmitting computer may transmit data to the computer, thus requiring
18 another ARP broadcast. In a network or branch with a large number of users, a great deal
19 of the network's bandwidth may be occupied by ARP broadcasting.

20 Consequently, it would be an advancement in the art to provide a method and
21 apparatus capable of reducing ARP broadcasting. The method and apparatus should enable
22 transmitting computers to obtain the MAC addresses of computers to which they will send
23 data, without propagating every ARP broadcast to every computer. Furthermore, the method
24 and apparatus should preferably reduce ARP broadcasting without the need to replace a great
25 deal of the currently-existing network infrastructure. The method and apparatus should be
26 inexpensive, low-maintenance, and fast. Finally, the method should be fully compliant with

1 existing protocols for network data transmission, so as to be transparent to computers and
2 end users on the network.

4 **BRIEF SUMMARY OF THE INVENTION**

5 In order to overcome many or all of the above-discussed problems, the present
6 invention comprises methods, apparatus, and systems for implementing Large-scale high
7 speed computer network. The network may connect an entire neighborhood or city in
8 networked communications, and accordingly, will be referred to herein as a Neighborhood
9 Area Network (NAN). The NAN of the present invention is a network conducted on a
10 unique scale with a unique clientele and is implemented in a manner that transcends
11 traditional network boundaries and protocols. The NAN is not equivalent to a wide area
12 network WAN, in part because it is essentially routerless. That is, while a plurality of NAN,
13 may be interconnected through the use of routers, each individual NAN is preferably
14 constructed without the use of internal routers. The NAN is unique from local area networks
15 (LANs) as well. One reason is that, due to its many novel features, it can be of a size and
16 scope previously unobtainable by conventional LANs.

17 The NAN is further unique because it is intended to cover and serve a selected
18 geographical area and to blanket that geographical area, rather than functioning to serve a
19 specific government, business, educational, or similarly related entity. Accordingly, the
20 subscribers and users of the NAN may be substantially non-related in any traditional business
21 manner. Furthermore, funding for the NAN, rather than being provided by a business-type
22 entity or subsidized by a governmental organization, may be funded at least in part by an
23 independent third party, such as a utility company and may be funded in total or in part by
24 subscribers.

25 The NAN is also comparatively inexpensive to install, making the placement of a
26 NAN in every neighborhood a real possibility. The NAN of the present invention is capable

1 of eliminating the message traffic burden from the Internet, thereby speeding up the Internet,
2 as it is adapted to be operated completely independent of the currently highly burdened
3 telecommunications infrastructure (although Internet service may be provided over the
4 NAN).

5 In one embodiment, the NAN is comprised of an optic fiber ring serving as the outer
6 backbone of the NAN. The ring is preferably populated with one or more fiber boxes, each
7 containing circuitry including switches, repeaters, gateways, etc. The fiber boxes in one
8 embodiment connect the backbone to a central office or headquarters data center in which
9 a server is preferably located. One or more gateways are preferably provided within the
10 backbone for access by Internet Service Providers (ISPs). An inner backbone comprised of
11 scalable 10 to 100 megabit coaxial cable preferably branches from the fiber backbone.

12 The coaxial cable preferably originates at the fiber boxes and branches through the
13 selected geographical region (discussed herein as a neighborhood, but of course, any
14 geographical scale could be served), connected by repeaters and nodes to individual
15 communicating stations. The inner backbone is preferably partitioned for efficient routing
16 of traffic.

17 The nodes in one embodiment comprise hubs. The repeaters may be placed three
18 hundred feet apart along the coaxial cable, with hubs placed within thirty feet of every house,
19 business, or other type of communicating station on the NAN. The hubs preferably connect
20 to the local houses or other buildings with ten-base-T twisted pair copper wiring employing
21 the Category 5 (Cat5) standard. The hubs in one embodiment are powered by one or more
22 of the communicating stations that they service. Accordingly, each station connected to a
23 hub may share the powering of the hub and may share the powering of other switching
24 equipment of the NAN as well.

25 In one embodiment NAN software operates on the server, the fiber boxes, the
26 repeaters, and the hubs. Client software preferably operates a computers located at each

1 communicating station. Additional functional software or logic may also execute on
2 communicating stations or computers of subscribing service providers. For example,
3 software may communicate with an electric power meter for transmitting information
4 regarding power consumption from a communicating station (the power customer) through
5 the network to third party service provider, in this case, a utility power company.

6 In one embodiment, at least a portion of the backbone is installed over the right-of-
7 way owned by or franchised to a public utility such as gas, electric, or power company. This
8 negates any need for a separate utility administering the NAN to acquire a new easement or
9 franchise from the landowners or the government entity of the geographic region. The NAN
10 may be financed and/or installed through the cooperation of the utility service provider
11 company. This arrangement allows the public utility service provider that would otherwise
12 be unable to enter the digital communication market to participate. It is also advantageous
13 in that a NAN developer or administration entity would otherwise likely be unable to afford
14 to finance and install the NAN due to the cost and risk of funding and lack of sufficient
15 rights-of-way.

16 In certain embodiments of an apparatus and method in accordance with the present
17 invention, an independent entity may create a city-wide network or NAN. The network
18 includes, in one embodiment, a fiber optic ring within the city to serve as a local backbone.
19 The fiber optic ring may be fully redundant. That is, it preferably completes a loop such that
20 any break in the loop will not shut the whole system down. The fiber can be laid
21 inexpensively as distances are not great and thus, less expensive local short-distance-types
22 of fiber cable can be used. A low cost fiber can be used, such as feeder fiber which is less
23 costly, and which requires less labor to install.

24 The fiber backbone is preferably populated by fiber boxes having switches therein.
25 Coaxial cable from switches to bridges and repeaters to hubs. The hubs may connect to
26 client stations using twisted-pair, copper cabling. A central server may be used and may be

1 located within a headquarters data center. A headquarters data center may be employed as
2 a gateway for Internet service providers. In addition, the Internet service providers may enter
3 the system through other gateways including one or more switches.

4 The fiber backbone may be laid using the franchise agreement granted to the power
5 company within a city or region. Thus, as the entire network is laid independently, the ISP
6 service is provided independent of the telecommunications line over the entire route.
7 Additionally, all ISPs are available on the net allowing equal access without choking traffic.

8 The infrastructure is preferably upgradable from 10 megabit to gigabit technology
9 over the same lines, such that the lines need not be relaid in order to upgrade. Services that
10 can be provided include surveillance, on-line books, two-way multi camera, schools, etc.
11 Additionally, IPBX, telephone, television, CATV, and video on demand can be provided
12 over the NAN. Video can be provided allowing independent selection, broadcast, start time
13 and may be buffered to the user in real time.

14 The NAN also preferably incorporates one or more multi-port switches which are
15 configured to truncate broadcast data. The multi-port switch is preferably an indoor switch
16 but is contained in an aluminum pedestal of dimensions approximately 3 by 2 by 2 feet and
17 is environmentally controlled.

18 The repeaters in preferred embodiments convert the data from the switches to be
19 transmitted over coaxial cable and are preferably semi-intelligent. In one embodiment, the
20 repeaters are housed out of doors within a protective pedestal. The pedestal may be located
21 on the ground or hung from power lines.

22 The bridges are, in preferred embodiments, high speed with a look-up binary tree and
23 are preferably contained in the protective pedestals. The bridges also filter out broadcast
24 traffic. The hubs route traffic to subscribing communicating stations and convert from
25 coaxial to twisted pair cable. The hubs are connected with a T-connector and powered by
26 the cooperative power coupler of the present invention.

1 The P-coupler preferably includes a series of transformers, one at each
2 communicating station. The communicating station connect with Cat5 wiring to the hub
3 through a home connection box. The home connection box preferably provides convenient
4 connections for power to the hub and for transmit and receive lines. The lines at the home
5 connection box are wired alphabetically. The home connection box connects preferably
6 connects with Ethernet cabling to a network card located within a computer at the client
7 station.

8 A modular power connector is preferably located at the home connection box. The
9 wiring from the communicating station to the hub operates, in one embodiment, at ten
10 megabytes per second. Three pairs of lines are preferably used, a transmit twisted pair, a
11 receive twisted pair, and an A/C twisted pair running from the transformer to power the hub.

12 The NAN of the present invention is a high speed routerless network which differs
13 from traditional large scale networks in that traffic is routed locally and that it has the speed
14 of a small local area network but with many more stations connected thereto. The large
15 amount of communicating stations is facilitated by the many novel aspects of the invention.

16 The NAN can be described as a baseband network rather than a broadband network
17 because it addresses communicating stations directly and linearly rather than through
18 broadcasting of data. The NAN of the present invention defines what cannot be routed rather
19 than defining the types of packets that can be routed. The NAN also preferably uses
20 converse/inverse filtering. Because the communications traffic is direct-routed, neighbor to
21 neighbor communication is very high speed and occupies only a small part of the NAN. It
22 also reduces the burden on the Internet.

23 Moreover, a packet prioritization method with an apparatus suitable for its
24 implementation is included to improve prioritization of packets leaving a switching station.
25 (A switching station refers to any device that performs switching between a plurality of ports,
26 regardless of whether the device is designated as a hub, bridge, switch, repeater, etc.) The

1 switching station may have a number of ports, each of which has a buffer to temporarily store
2 incoming packets. The switching station may also have a processor and program memory
3 containing instructions for the processor. A cache may also be provided for additional data
4 storage, with a multiplexer to enable the cache to simultaneously receive signals from
5 multiple sources. The processor, buffers, and multiplexer may all be linked by a bus.

6 Similarly, a packet prioritization station is provided, either as an integral part of the
7 switching station, or as an addition, such as an auxiliary expansion card or board (AEC). If
8 embodied as an AEC, the packet prioritization station may have a bus linked to the bus of
9 the switching station by an interrupt controller that triggers the packet prioritization station
10 when the proper conditions are met in the switching station. The switching station, in its
11 independent form, has a processor and a program memory, both of which may take multiple
12 forms. The processor carries out instructions provided by the program memory in order to
13 carry out the functions of the packet prioritization station.

14 A cache in the packet prioritization station contains a database binding each MAC
15 layer address (or destination) to other MAC layer addresses (or origins) that have sent
16 packets to that MAC layer address. These destinations and origins are obtained by copying
17 them from a sampling of all packets passing through the switching station. The origins are
18 maintained in the database for a certain period of time. The processor, program memory, and
19 cache of the packet prioritization station are all linked by the bus.

20 When a new packet is received through an incoming port of the switching station, it
21 is stored in the buffer for the incoming port. Meanwhile, the switching station matches it up
22 with one or more ports, through which it will be transmitted to reach its destination. When
23 packets in multiple buffers are not routed to a single outgoing port, the packets in the buffers
24 are simply sent to their respective outgoing ports in cyclical, or "round robin" fashion.
25 However, when more than one packet is routed to an outgoing port, the packet prioritization
26 station must determine which packet get priority.

1 It has been discovered that those packets being sent to destinations for which many
2 origins are cached typically are of a higher relative importance, because they represent
3 multiple users or time-intensive network use. The packet prioritization station proceeds
4 through packets routed to a single port in round robin format, until it encounters the first
5 packet with a destination having more than a threshold number of origins bound to it. That
6 packet is immediately sent. When no packet routed to the outgoing port has a destination
7 that has recently received packets from the threshold number of origins, packets are sent in
8 round robin fashion, i.e., by sending packets from alternating incoming ports. The process
9 continues until traffic to that outgoing port subsides.

10 Consequently, destinations receiving data from many sources will receive priority.
11 More time-critical communications are transferred first, because smaller files, such as e-mail,
12 are typically those for which rapid response is especially important. Large information
13 transfers, such as file downloads, normally are not as critical, and can therefore be delayed
14 until after more important information has been routed. Similarly, files from multiple origins
15 are often sent to multiple recipients. Thus, the packet prioritization station handles the needs
16 of the majority of users as rapidly as possible.

17 Furthermore, a traffic reduction method and apparatus may also be implemented
18 according to the present invention. An ARP caching station may be provided to work in
19 concert with the switching station. The ARP caching station may be used with or without
20 the packet prioritization station, and the packet prioritization station may similarly function
21 independent of the ARP caching station.

22 The ARP caching station may also be integral with the switching station, and may
23 share its components for operation. Alternatively, the ARP caching station may be an AEC
24 with independent componentry, in communication with the switching station. Thus, the ARP
25 caching station may have its own processor, program memory, and cache, linked by a bus.
26 As with the packet prioritization station, the bus of the ARP caching station may be linked

1 to the bus of the switching station by an interrupt controller. Thus, operation of the ARP
2 caching station may also be triggered by the switching station.

3 The ARP caching station may have its own database containing associated IP
4 addresses and MAC addresses. These may be obtained by storing the addresses from any
5 packet, such as an ARP broadcast response, that contains both an IP address and a MAC
6 address denoting the same destination.

7 When an ARP broadcast is received by the switching station, the ARP caching station
8 may be activated to look for the designated IP address in the cache, and return the associated,
9 requested MAC address if it is available. If the requested MAC address is not found in the
10 cache, the ARP caching station may store the designated IP address in the cache for future
11 reference, or may simply store nothing until another packet with an IP address and a
12 matching MAC address is received. In any case, the ARP broadcast is then propagated by
13 the switching station so that a response can be sent by the destination (the computer having
14 the designated IP address).

15 If the requested MAC address is available, the ARP caching station creates a packet
16 in the proper form for an ARP broadcast response containing the requested MAC address,
17 and sends it to the originator of the ARP broadcast. The ARP broadcast need not be
18 propagated by the switching station. Thus, if the requested MAC address is in the cache,
19 considerable bandwidth is saved by avoiding transmission of the broadcast through all ports
20 on the switch (except the incoming port of the ARP broadcast). The originator of the
21 broadcast also receives a quicker response and can begin transmitting information to the
22 destination with little delay.

23 24 **BRIEF DESCRIPTION OF THE DRAWINGS**

25 In order that the manner in which the above-recited and other advantages and objects
26 of the invention are obtained will be readily understood, a more particular description of the

1 invention briefly described above will be rendered by reference to specific embodiments
2 thereof which are illustrated in the appended drawings. Understanding that these drawings
3 depict only typical embodiments of the invention and are not therefore to be considered to
4 be limiting of its scope, the invention will be described and explained with additional
5 specificity and detail through the use of the accompanying drawings in which:

6 Figure 1 is a schematic block diagram illustrating one embodiment of network system
7 hardware for use with the present invention.

8 Figure 2 is a schematic block diagram illustrating one embodiment of a system
9 architecture for use with the present invention.

10 Figure 3 is a schematic block diagram of one embodiment of a network architecture
11 for use with the present invention.

12 Figure 4 is a schematic block diagram of one embodiment of a traffic filter module
13 for use with the present invention.

14 Figure 4A is a schematic representation of one embodiment of a communications
15 packet of the present invention.

16 Figure 4B is a schematic representation of an OSI seven layer model.

17 Figure 5 is a schematic block diagram of one embodiment of a switching station, a
18 packet prioritization station, and an ARP caching station suitable for use in the present
19 invention, linked by interrupt controllers.

20 Figure 6 is a schematic block diagram of a buffer suitable for use in the switching
21 station of Figure 5, in which incoming packets are stored.

22 Figure 7 is a schematic block diagram of a program memory suitable for use in the
23 switching station of Figure 5, with various executable modules to carry out the functions of
24 the switching station.

25 Figure 8 is a schematic block diagram of a cache suitable for use in the switching
26 station of Figure 5, with a table of MAC addresses and associated ports.

Figure 9 is a schematic block diagram of a program memory suitable for use in the packet prioritization station of Figure 5, with executable modules to carry out the functions of the packet prioritization station.

Figure 10 is a schematic block diagram of a cache suitable for use in the packet prioritization station of Figure 5, with a table of destination MAC addresses, each of which is associated with one or more origin MAC addresses.

Figure 11 is a schematic block diagram of a program memory suitable for use in the ARP caching station of Figure 5, with executable modules to carry out the functions of the ARP caching station.

Figure 12 is a schematic block diagram of a cache suitable for use in the ARP caching station of Figure 5, with a table of IP addresses associated with MAC addresses.

Figure 13 is a flowchart diagram of a method suitable for carrying out the invention, in which a packet is received and processed by a switching station, packet prioritization station, and ARP caching station.

Figure 14 is a flowchart diagram of a packet receiving step suitable for the method of Figure 13.

Figure 15 is a flowchart diagram of a priority information storage step suitable for the method of Figure 13.

Figure 16 is a flowchart diagram of an ARP request processing step suitable for the method of Figure 13.

Figure 17 is a flowchart diagram of an address caching step suitable for the method of Figure 13.

Figure 18 is a flowchart diagram of a packet routing step suitable for the method of Figure 13.

Figure 19 is a flowchart diagram of a blocking decision step suitable for the method of Figure 13.

1 **DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS**

2 Referring now to Figure 1, shown therein is a schematic block diagram showing
3 various hardware components of one embodiment of a large-scale, high speed network of the
4 present invention. Because the network is intended to serve a selected geographical region,
5 it is referred to herein as a neighborhood area network ^{NAN} (ANA) 10. The NAN 10, as depicted,
6 includes a backbone 12, that is divided into two components. A first component is a fiber
7 backbone 14 that is preferably adapted to transmit packetized data using standard optical
8 communications protocols and technology. The fiber backbone 14 is preferably configured
9 in a ring with incoming traffic traveling in a selected given direction.

10 A second component comprises a local backbone 16 that is preferably configured
11 with a non-redundant branching structure and that is adapted to transmit data using radio
12 wave signals. In the schematic depiction of Figure 1, the physical locations of connections
13 are represented, while an example of the actual branching structure is shown in Figure 3.

14 The NAN system 10 in the depicted embodiment of Figure 1 also includes a server
15 18 which may be located at a central headquarters office 20. One or more fiber switches 22
16 may be located within the fiber backbone 14. Indeed, the fiber backbone 14 may complete
17 a circle around a neighborhood or other common geographical region which is intended to
18 be networked in computer, voice, and or/video communication. The fiber backbone 14 may
19 be provided with redundant loops in case one loop becomes inoperable.

20 The local backbone 16 preferably communicates with the fiber backbone 14 through
21 one or more fiber switches 22. Each fiber switch 22 is preferably configured to examine
22 packetized message traffic passing therethrough, and where a message is intended for a
23 communicating station serviced by a portion of the local backbone serviced by the switch 22,
24 route the message onto the local backbone 16. Each switch 22 also preferably routes locally
25 generated traffic with external destinations to the fiber backbone 14 for receipt by other
26

1 switches or gateways 108 to the Internet 34. The switches 22 preferably also convert
2 communications between optical communications signals and radio frequency signals.

3 Within the local backbone 16, switching devices, including a series of repeaters 24,
4 nodes 26, and bridges 50 are preferably deployed. In one embodiment, the local backbone
5 16 is provided with coaxial cable 38 having a sufficiently high band width and having signals
6 of sufficiently high amplitude that repeaters 24 are needed only every 300 feet or so. The
7 nodes may comprise hubs 26 which, due to the efficient propagation of the NAN 10, can be
8 located up to 30 feet from each communicating station 30.

9 Communicating stations 30 in one embodiment connected to the nodes 26, with Cat
10 5, twisted pair wiring 40 through a home connection box 42. Internet Service Providers
11 (ISPs) 32 are shown connected to the NAN 10 through in several different types of gateways.
12 An ISP 32 may connect through the central headquarters office 20 and from there to a fiber
13 switch 22. Alternatively, an ISP may communicate directly with the fiber backbone 14
14 through a fiber switch 22. The ISPs provide access to the worldwide web and the Internet
15 34.

16 Each communicating station 30 may be provided with one or more home service
17 boxes 44. The service boxes 44 communicate over the NAN 10 and provide interactivity
18 from a remote distance. The service boxes 44 may comprise, for instance, power meters 46,
19 security systems 48, and any number of electrical and mechanized devices, including
20 appliances, sprinkling systems, synchronized clocks, etc.

21 The fiber switches 22 may be housed within containment units 52. The containment
22 units 52 may be located inside or out of doors and are preferably provided with insulation
23 and/or environmental control devices such as a fan 54 and/or air conditioning 56. The
24 containment units 52 are preferably vented.

25 The repeaters 24, bridges 50 and nodes 26 are preferably located within protective
26 pedestals 28 which are also preferably vented, which provide a hardened outer shell, and

1 which may be provided with fans 54 or other environmental control devices. The pedestals
2 28 may be mounted in the ground, or may be mounted from utility and/or power lines
3 overhead. The pedestals 28 preferably provide some type of lightening protection such as
4 a Faraday shield. The pedestals 28 are described in greater detail below with reference to
5 Figures 7 and 8.

6 Figure 2 is a functional block diagram illustrating a system architecture 100 including
7 operative data structures and executable modules for controlling the operation of the
8 hardware of the NAN 10 depicted in Figure 1. The system architecture 100 controls the
9 interactions of the various intelligent components of the NAN 10 of Figure 1.

10 Accordingly, shown in Figure 2 are the different modules and executables for
11 operating the NAN 10. Included are a plurality of client stations 30 communicating over a
12 transmission system 102. Other entities may also communicate over the transmission system
13 102. These include the central headquarters office 20, the server 18, a monitoring station
14 152, and service providers 104, including a utility company 106.

15 Referring now to the transmission system 102, one method of operation of the NAN
16 10 to transmit information between the client stations 30 will be described. In one
17 embodiment, the NAN backbone 12 is essentially routerless. That is, the system is operated
18 at a large scale, but using the same principles as a small local area network. This is
19 achievable due to the unique architecture and configuration of the NAN 10. Routers (62 in
20 Figure 3) are required only when connecting to outside entities, such as other NANs or the
21 Internet 34.

22 Components included within the system 100 include the bridges 50, the switches 22,
23 the repeaters 24, and the nodes, which in one embodiment comprise hubs 26. Also included
24 within the system 102 is an Internet routing module 108 which routes traffic to and from the
25 ISP's 32. The Internet routing module 108 operates as a gateway and may comprise a switch
26 22 and a router 62.

1 The switches 22 are provided with software modules in the form of a switch routing
2 module 110 and a switch conversion module 112. The switch routing module 110 is used
3 to route traffic between the switches 22. The switch conversion module 112 is used to
4 convert packeted traffic between the optical communications protocol and the radio
5 frequency signals used within the coaxial cable lines 16. Thus, in preferred embodiments,
6 each switch includes one or more protocol converters interfacing between fiber cabling and
7 Cat5 twisted pair wiring.

8 The protocol converters translate the optical signals into radio frequency signals for
9 transmission on the coaxial Cat5 cables. The radio frequency signals are in turn translated
10 into digital signals by the network cards 156.

11 The Cat5 twisted pair wires lead into out of the switch 22 and connect to the protocol
12 converters 112 and to repeaters 24. The repeaters 24 place the data packets on the coaxial
13 cable 16. The Cat5 wiring may also lead directly to client stations 30 that are within 300 feet
14 of the switch 22.

15 Traffic is routed in an efficient manner whereby the system 100 utilizes the high
16 speed fiber cables 14 to as great a degree as possible routing packetized traffic to the switch
17 22 closest to the communicating station 30 to which the message is addressed. Once the
18 packet reaches the closest switch 22, it is routed through a repeater 24 onto the local
19 backbone 12. Once on the local backbone 12, the packet passes to a bridge 50 and then to
20 the node 26 closest to the client station 30 in a manner be discussed below with relation to
21 Figure 3.

22 The repeaters 24 are preferably spaced approximately every 300 feet in order to avoid
23 over-attenuation of the signals carrying the data packets. The nodes 26 are placed within 30
24 feet of each communicating station 30.

25 The communicating stations 30 are preferably provided with client software 126 for
26 enabling communications over the NAN 10. The NAN 10 communications medium is, in

1 one embodiment, standardized Ethernet data packets adhering to the Ethernet/OSI standards.
2 In one embodiment, the data packets may be transmitted over the NAN 10 using merely
3 MAC addresses of the low levels of the OSI model.

4 Client stations 30 which are new to the NAN 10 transmit an initial communication
5 packet over the NAN 10 to the server 18. The server 18 in reply issues an IP address 136 to
6 the client station 30 which is semi-permanent. Thereafter, the client station 30 has a semi-
7 permanent IP address 136 which is changed only upon incidents such as the computer or
8 network card of the client station 30 being changed.

9 The packets are routed through the switches 22, repeaters 24, and nodes 26, to the
10 addressed client stations 30. The packets may be transmitted at a rate of 10 megabits per
11 second due to the unique architecture of the NAN 10. This high rate of speed can be
12 upgraded by a factor of 10 or even up to a factor of one hundred without having to redeploy
13 the fiber cables 14, the coaxial cables 16, and the pair twisted wiring 40. This, again, is due
14 to the unique architecture of the system.

15 The system architecture includes extending the distance a packet can travel up to
16 between 3000 and 25000 feet and increasing the maximum tolerable packet acknowledgment
17 time. This is accomplished in one embodiment by digressing from the IEEE standards.

18 For instance, the signals with which the packets are transmitted are amplified to a
19 higher power than those on standard networks. This is accomplished by increasing the gain
20 in the amplifiers that make the repeaters function. Additionally, the reception equipment is
21 preferably more sensitive and able to capture a more degraded signal than standard network
22 equipment.

23 The fact that the system operates on a baseband concept wherein all of the cable
24 bandwidth is restricted to one channel rather than being divided into multiple channels
25 allows for a higher bandwidth and greater power from the repeaters. This allows for collision
26 detection over the cable 38 and for a release of the collision detection at a much lower level.

1 Thus, voltage spikes are detected and ignored so that lower level collisions are not detected
2 and the large level collisions can be detected. The incidences of these collisions are highly
3 reduced due to the high bandwidth and direct routing of the system 100.

4 Collision detection is preferably accomplished through voltage detection and timed
5 resends and is adjusted to compensate for the increased sensitivity of the repeaters.

6 The repeaters 24 are provided with software or other logical circuitry 120 therein
7 which allows the repeaters 24 to be semi-intelligent. The repeaters 24 transmit the fact that
8 they are functioning, as well as information regarding the amount of traffic passing
9 therethrough, in order to better manage the NAN 10. Otherwise, the repeaters 24 merely pass
10 the packets through and do not provide any switching function, merely increasing the
11 amplitude of the signals carrying the packets. As mentioned, the repeaters 24 are, in one
12 embodiment, placed every 300 feet across the local backbone 16.

13 The hubs 26 route the packetized traffic through the Cat5 twisted pair wiring 38 to
14 the communicating stations 30. Internet routing 108 may also take place to route the Internet
15 communications to the ISPs 32. Communications with external stations over the Internet 34
16 may be conducted with a permanent IP address to get the messages within the NAN 10,
17 wherein the outside data packets are routed using MAC addresses. Additionally, stations 30
18 without permanent IP addresses may communicate through the use of a masqueraded IP
19 address using a permanent IP address to get into the NAN and the semi-permanent IP
20 addresses 136 issued to each client station 30 in a manner that will be discussed below in
21 greater detail.

22 The bridges 50 are provided with software 114 and are also provided with a memory
23 116 containing a bank 118 of the IP addresses 136 of each client station 30. The bank 118
24 also includes, for each corresponding IP address 136, information regarding the location of
25 the client station 30 to which the IP address 136 is assigned.

26

1 Accordingly, the bridges limit communications to only a particular portion of the
2 network 10 to which the communication is addressed. Thus, the bridges 50 effectively
3 partition the NAN 10. A further function of the bridges 50 and the switches 22 is to
4 eliminate unwanted communications. For instance, in one embodiment, broadcast packets
5 and messages are forbidden. Accordingly, each switch 22 and bridge 50 may be provided
6 with a traffic filter module 160 as depicted in Figure 4.

7 Referring to Figure 4, the traffic filter module 160 is used to eliminate certain types
8 of traffic that may not be routed over the NAN 10. Accordingly, the NAN 10 is defined as
9 determining what types of communications can not be routed rather than determining what
10 types can be routed, as in the prior art. Within each traffic filter module 160 may be a
11 broadcast traffic sniffing module 162. The broadcast traffic sniffing module 162 examines
12 each information packet 165 (shown in Figure 4A) and checks certain fields 171 which
13 indicate that the packet 165 is broadcast data. When the traffic sniffing module 162
14 determines that the packet 165 is broadcast traffic, it then initiates the traffic elimination
15 module 164 which eliminates the broadcast packet 165.

16 The bridges 50 and switches 22 in one embodiment detect broadcast traffic by
17 detecting an empty field 171 within the MAC address 170. Alternatively, the broadcast
18 traffic sniffing module 162 may detect a series of addresses at a certain level such as 255,
19 255, 255 to detect a broadcast packet 165.

20 Thus, because the NAN 10 eliminates unwanted traffic and restricts traffic to only
21 those portions of the NAN 10 through which the packet 165 must travel to reach the
22 addressed communication station 30 in the most efficient manner, much extraneous traffic
23 is eliminated. This, combined with the higher speeds of the present invention, allow the
24 NAN 10 to be operated as if it were a local area network but on much grander scales, indeed,
25 even to include entire neighborhoods or municipalities. Additionally, because of this, the
26 NAN 10 is suitable for use in geographical areas covering extensive distances that are merely

1 geographically or community interest related, rather than being business, government,
2 education or otherwise related. Thus, the NAN system 10 can be by financed at least in part
3 by the service providers which will benefit from the efficient communication of the NAN 10.

4 Referring now to the service providers 104 of Figure 2, an example of such a service
5 provider is a utility company 106. In one embodiment, the utility company 106 is a power
6 company. Thus, for example, the power company can communicate over the transmission
7 system 102 on the NAN 10 with each client station 30. Within each client station 30 is one
8 or more service boxes 144 having therein customer service software 150.

9 The customer service software 150 might, in one instance, comprise power meter
10 software 148 within a power meter box 46. The power meter software 148 may transmit
11 power usage through the NAN 10 back to the utility company 106. The utility company 106,
12 with a power usage collection module 144, receives the power usage data and transmits it to
13 a billing module 146. The billing module 146 then bills the customer at the communicating
14 station 30 over the transmission station 102. The payment of the bill may also pass through
15 the transmission system 102, thus passing through the NAN 10 back to the utility company
16 106. Of course, utility companies other than the power company may also use this system
17 of data collection billing and payment receipt.

18 Other types of service boxes 144 may also contain customer service box software
19 150. For instance, the security system 48 may contain therein software which notifies the
20 monitoring station 152 of any irregularities. Software 154 within the monitoring station 152
21 may monitor the data transmitted by the security system 48. For instance, this data might
22 include home security system data indicating that a break-in has occurred. The security
23 system 48 may also indicate the occurrence of a fire, and may transmit full video surveillance
24 data back to the monitoring station 152. The monitoring station 152 or a similar station may
25 also monitor the contents of the NAN 10 in order to eliminate illegal traffic. Pornography
26 or other types of traffic may likewise be eliminated.

1 Each client station 30 as mentioned, preferably communicates at the MAC layer
2 within the NAN 10. The client stations 30 may also be provided with a semi-permanent IP
3 address for communications external to the NAN 10. The server 18 is provided with server
4 software 124 which maintains a bank 138 of the IP addresses 136. The server 18 thus issues
5 the IP addresses 136 and also maintains a binding between the MAC layer communications
6 and the IP addresses 136. These bindings are transmitted to the switches 22, bridges 50, and
7 any other equipment with a need to know the IP addresses 136 of the client stations 30.

8 Consequently, the server 18 is not necessary other than for issuing IP addresses and
9 maintaining bindings, and indeed, if the server 18 were to go down, the transmission system
10 102 operating on the NAN 10 could continue to operate. New client stations 30 would
11 merely not be able to receive an IP address.

12 The central headquarters office 20 preferably contains therein a headquarters software
13 module 128. The headquarters software module 128 may conduct monitoring and billing
14 types of operations. Thus, a customer database 130 may be maintained therein and may
15 coordinate with a billing module 134. A redundant database 132 is also preferably included.
16 The redundant database 132 may be located at a distant site such that it maintains a copy of
17 all data in the case of a failure of the customer data 130. Synchronizing information may
18 pass between the customer database 130 and the redundant database 132 over the NAN 10
19 with the use of the transmission system 102.

20 Billing information may be generated and stored within the billing module 134 and
21 may be transmitted to communicating stations 30 over the transmission system 102. The
22 customer database 130 may maintain records including records of which customers are
23 behind on their payments. If the customers are behind, the client station 130 of that customer
24 may be denied services in part or in full of the NAN system 10. These services include, in
25 one embodiment, Internet service.

26

1 The communicating stations 30 are preferably provided with standard network cards
2 156 which transmit through the home connection box 42. The client software 126 residing
3 at the communicating stations 30 preferably maintains the client's IP address 136 and
4 receives and generates data packets (shown at 165 in Figure 4A) with which information is
5 transmitted over the transmission system 102. The client software 126 may provide many
6 various types of functions, including video phone communication, audio, and video
7 transmission, payment of bills, ordering of on-demand video, transmission of home security
8 information, etc.

9 A power coupler 135 may be provided within or in communication with the home
10 connection box 42. The power coupler 135 preferably conditions incoming power from a
11 power source at each communicating station, combines the power and network connection,
12 and provides a simple manner of connecting the twisted pair wiring to standard computer
13 cabling, preferably Ethernet cable, which passes to the computer at the communicating
14 station 30. In one embodiment, the twisted pair wiring is provided with a twisted pair for
15 transmission, a twisted pair for reception, and a twisted pair carrying AC to the hub 26, as
16 will be discussed in greater detail below with reference to Figures 5 and 6.

17 The hub 26 is in one embodiment provided with a power concentrator 25 which
18 provides power conditioning and power delivery to the hub 26. The power concentrator
19 receives power from the power coupler 135 of the communicating stations 30. Preferably
20 the power concentrator 25 receives power from two or more stations 30 and passes the power
21 on to the hub 26 or other switching device. A Power concentrator 25 receives power through
22 a transformer connected to a wall socket at the communicating station 30. In one preferred
23 embodiment, four houses share a hub and provide power to the hub. The hub bleeds power
24 out of the four transformers at a time, but can receive power from less than all of them and
25 be at a full power level. This redundant power supply scheme ensures that the hub 26
26 continues operating even if one of the power sources, i.e., one of the communicating station

30, goes down. Thus, AC power is received from the communicating station 30 through the power coupler 135 to the power concentrator 25. In addition, all switching equipment may be powered cooperatively in this manner and may be provided with power concentrators 25.

In one embodiment, the AC power is received directly from a power meter at the communicating station 30. The power from the communicating stations 30 may be provided individually or collectively to the switches, bridges, repeaters, router, hubs, and any other switching equipment of the NAN. Additionally, power meters not located at communicating stations 30 may be utilized to provide power to the hubs 26 and other switching equipment.

In one embodiment, the communicating stations 30 or the hubs 26 comprise a power meter monitoring hub 26. The power meter monitoring hub 26 may comprise an RF receiver and an 8-bit microcontroller as well as an RS 232 communications interface and a power supply. The hub may also contain up to four 10-base T ports. On-site configuration is provided by an RS 232 port. Under this embodiment, the monitoring hub receives power consumption data from power meter transmitters and passes it on to the utility company 106 over the transmission system 102.

Each power meter 46 in this embodiment provided with a power monitoring transmitter. The transmitter may be comprised of a PIC microcontroller, a 418 megahertz UHF transmitter, a photo-reflective sensor, and an off-line power supply. The transmitter may use the photo-reflective sensor to monitor rotation of the power meter disk and store the information in nonvolatile memory in the microcontroller. The transmitter transmits the power usage information to the power meter monitoring hub along a 418 megahertz RF link.

In one embodiment, the coaxial cable, as well as the 10-base T wire, is housed within a protective conduit. The system may operate with Linux using an IP chain and masquerading which is considered more effective than using a proxy server.

The bridges 50, in addition to eliminating broadcast traffic, may also receive and regenerate the packets 165 at a higher power level. The repeaters 24 preferably merely

1 amplify the signals carrying the packets 165 and do so without any delay, while the bridges
2 may slow down the packets somewhat.

3 Referring now to Fig. 3, shown therein is a functional block diagram of a NAN
4 hierarchy scheme 60. Within the scheme 60 is shown the fiber backbone 14 looping in a
5 circuitous manner to form a ring. Within the fiber backbone 14 is a plurality of switches 22.
6 A central switch 22a is shown connected with the central headquarters 20 and through a
7 router 62 to the Internet. Thus, the fiber backbone 14 comprises an outer circuitous
8 backbone. It should be noted that the NAN 10 may have a plurality of gateways 62. Because
9 of the plurality of gateways, any number of ISP providers 32 may provide service to the NAN
10 10. Other types of service providers and outside entities may also access the NAN 10
11 through the gateways 62.

12 Emanating from the switches 22 are components of the local backbone 16 which are
13 arranged in a branched configuration. Thus, shown branching out from each switch 22 is a
14 series of bridges 50, repeaters 24, and hubs 26. Each bridge 50 separates and services a
15 plurality of hubs 26.

16 Thus, an incoming packet 165 received, for instance over the Internet 34, passes
17 through the router 62. The router 62 uses an IP address 169 shown in Fig. 4a to determine
18 that the packet is local to the NAN 10. For instance, the IP address may be assigned to the
19 NAN 10 or to the router 62 specifically under a masquerade scheme that will be described.

20 Once the packet 165 reaches the NAN 10, it is routed using a MAC address 170 of
21 Fig. 4a. After passing through the router 62, the packet 165 is received by the central switch
22 22a. As shown in Fig. 4A, the packet 165 comprises a header 166, a data portion 167, and
23 a footer 168. The header comprises the address of the addressed communicating station 30.
24 The footer contains redundancy information to make sure the packet 165 was properly
25 received. A cyclical redundancy check (CRC) may be used using information in the footer
26 for acknowledgment that the packet 165 was received and has not been degraded.

1 Within the header 166 may be both an IP address 169 and a MAC address 170. The
2 MAC address 170 refers to a unique number given to each network card 156 of Figs. 2 and
3 5. The IP addresses 169 are administered by the Internic agency and are addresses utilized
4 under the TCP/IP protocol. Each station has a unique MAC address. Additionally, each
5 station may have a unique IP address 169.

6 Nevertheless, because IP addresses 169 are becoming scarce and difficult to procure,
7 a masqueraded system may be employed wherein the router 62 contains a routable IP
8 address or several routable IP addresses and stations 30 within the NAN 10 are addressed by
9 the routable IP address of the router 62 outside the NAN 10. Once addresses containing the
10 masqueraded IP address reach the NAN 10 at the switch 22a, the MAC address 170 may then
11 be used to route the packet 165 within the NAN 10. Indeed, within the NAN 10, routing is
12 preferably exclusively conducted using the MAC address 170.

13 When communicating on the MAC level, a communicating station 30, in one
14 embodiment, uses a protocol such as an ARP request. The "ARP" request is an address
15 resolution protocol. The ARP protocol talks to the network cards looking for the MAC
16 address. The use of an ARP-type address protocol by the NAN 10 does not adhere exactly
17 to the ARP address protocol but is similar to it.

18 Thus, the server 18 may be characterized as a modified DHCP server but does not
19 broadcast DHCP as with the prior art systems, though it does maintain the IP-MAC address
20 binding and notifies all subscribing components of that binding. Under this arrangement,
21 when a communicating station 30 comes on-line and receives the non-routable IP address
22 from the server 18, it then binds the IP address. In one embodiment, this is done by
23 populating its registry with the IP address. That is, the IP address is bound to the TCP/IP
24 protocol stack. This IP address is used for TCP/IP protocol communications with stations
25 72 external to the NAN 10. As discussed, all internal communications are preferably routed
26 using the MAC address.

1 Of course, the communicating stations 30 could also receive permanent IP addresses
2 either from the server 18 or directly from Internic. These permanent, routable IP addresses
3 may also be maintained within the binding of the server 18.

4 Preferably, hubs, bridges and switches work on only the lower two levels of the OSI
5 model of Fig. 4b. When a packet 165 is addressed to go outside of the NAN-10, it is sent to
6 the router 62 which acts as a gateway to the Internet 34 and passes the packet 165 outside the
7 NAN 10. The IP addresses within the communicating stations 30 communicate through
8 virtual ports on the communicating stations 30 but preferably not through the same
9 communicating ports as traditional DHCP protocol standards.

10 Additionally, the IP addresses are semi-permanent. That is, the communicating
11 stations 30 maintain a single IP address for external communications and do not flood the
12 NAN 10 with requests for DHCP servers to receive IP addresses from. Indeed, because of
13 this substantially, only direct routed traffic exists on the neighborhood, and all broadcast
14 traffic is substantially squelched. Additionally, all traffic is partitioned within its own area
15 and does not travel across the entire network. For this reason, there are substantially less
16 collisions because traffic is much more localized. This also allows the network to service
17 many more communicating stations 30.

18 The OSI model 190 is shown in Fig. 4b. As shown therein, the OSI model comprises
19 a first layer 191 known as the physical layer. A second layer 192 is known as the data link
20 layer and it is this layer that predominantly deals with the MAC address 170. A third layer
21 193 is referred to as the network layer, a fourth layer 194 is referred to as a transport layer,
22 and a fifth layer 195 is referred to as a session layer. The session layer 195 primarily deals
23 with the IP address 169. A sixth layer 196 is referred to as the presentation layer, and a
24 seventh layer 197 is referred to as the application layer. Within the seven layer OSI model,
25 the upper levels allow two communicating stations, one assigned as a client and one assigned
26 as a server, to coordinate communications with each other.

1 The NAN 10 may be configured to communicate only on the second layer 192 within
2 the loop of the fiber backbone 14. For example, the router 62 may be configured to receive
3 IP addresses 169 from the Internet 34, and provide only MAC addresses 170 to the switch
4 22. IP address resolution may be handled by the ISP 32 or other suitable entity. Thus,
5 communications between the communication stations 30 would occur using only MAC
6 addresses 170, without the need to send an IP address 169 within each packet 165.

7 The bridges 30 may then be omitted, and the hubs 26 and repeaters 24 may be
8 replaced by switches 22 configured to handle only the MAC addresses 170. Such an
9 architecture would provide more rapid data transmission throughout the NAN 10, since there
10 is less information in each packet 165 to deal with. In addition, installation and
11 configuration of the NAN 10 would be simpler because switches 22 may be installed without
12 any need for hubs 26, repeaters 24, and bridges 30.

13 Referring back to Fig. 3, once message traffic 165 is received from the router 62 to
14 the switch 22a, the switch 22a maintains the packet 165 momentarily in a buffer 164 and
15 refers to a database 66 to determine whether the MAC address 170 is local to a partition 169
16 belonging to the switch 22a. Switch 22a makes this binary determination, and if the answer
17 is yes, passes the packet 165 to a first bridge 50a.

18 If the answer is no, that is, the traffic is not local to a partition 168, the switch passes
19 the packet 165 in a given direction to a subsequent switch 22. In the depicted embodiment,
20 the given direction is clockwise. Upon passing the packet 165 on, a subsequent switch 22
21 receives the packet 165 and similarly examines the packet 165 to determine whether it is
22 local or external to a partition 168. If the packet is local to the partition 168, the switch 22
23 will pass it on to a bridge 50 within a partition 168 to which the switch 22 belongs. If the
24 packet 165 is addressed external to the partition 168 of the switch 22, the switch 22 passes
25 the packet 165 in the given (clockwise) direction to a subsequent switch 22.

26

1 Presuming that the packet 165 was local to switch 22a, switch 22a passes the packet
2 to a first bridge 50a. The bridge 50a then holds the packet 165 temporarily in a buffer 64 and
3 refers to a local database 66 to determine whether the packet 165 is local or external to the
4 bridge 50a. If the packet 165 is local to the bridge 50a, the bridge 50a determines which of
5 the hubs 26 connected with the bridge 50a the packet 165 must be routed through.

6 If the packet 165 is addressed external to the bridge 50a, the bridge 50a passes it to
7 a subsequent bridge 50b. The bridge 50b then receives the packet 165 within a buffer 64 and
8 examines its database 66 to determine if it the packet is addressed to a local station 30. If it
9 is not, it passes it on to subsequent bridges 50 (not shown) in the branching structure of the
10 local backbone 16.

11 The bridges 50 are typically separated by one or more repeaters 24 to amplify the
12 radio frequency (RF) signals which contain the packets 165. Referring now back to bridge
13 50a, if the packet 165 was local to bridge 50a, it determines which of the hubs 26 to pass it
14 to. Presuming that the packet 165 was addressed to a station 30a within a hub 26a, the bridge
15 passes the packet to the hub 26a. The hub 26a briefly maintains the packet 165 within a
16 buffer 64 and examines its database 66 to determine which of the subscribing communicating
17 stations 30 the packet 165 belongs to. In this case, it determines that the packet belongs to
18 station 30a and places the packet on a line 40 to be received by a network card 156 located
19 at the communicating station 30a. A similar process would occur with every bridge 50.
20 Thus, for instance, if the packet were addressed to a station 30b, the bridge 50b would
21 receive the packet and transmit to the hub 26b, which would receive the packet 165 and
22 transmit it to the communicating station 30b.

23 Inter-NAN communications are even more simplified. For instance, if the
24 communicating station 30a wishes to communicate with the communicating station 30b,
25 client software 126 would prepare the packet 165 and place it through the network card 156
26 onto the NAN 10. The packet 165 would be received by hub 26a which would in turn

1 transmit the packet 165 to the bridge 50a. The bridge 50a would examine the packet once
2 again to determine whether it is local or external to the bridge 50a. If it is locally addressed,
3 the bridge 50a transmits to the appropriate hub 26 connected thereto. If it is not, it directs
4 the packet 165 to another bridge 50 or to the switch 22a, depending on the MAC address 170.

5 The switching equipment, such as the switches, bridges, and hubs, preferably use a
6 binary tree sorting algorithm to sort through addresses in the attendant databases 66 to
7 determine the location of stations 30 addressed by the packets 165, which greatly enhances
8 the speed thereof. The binary tree, rather than being just a one dimensional look-up table or
9 bubble sort, is branched and allows for larger databases without significant propagation
10 delays. The binary tree is implemented, in one embodiment, using the Nikolas Wirth style
11 that is known in the art.

12 Note that each bridge 50 also preferably contains its own sub-partition 70 in the
13 partition 68 of the switch 22 to which it subscribes. In this case, when a bridge, such as
14 bridge 50 determines that the packet 165 is local to the partition 68 but not within its own
15 subscribing hubs 26, the bridge 50a passes the packet 165 on to the bridge, e.g. bridge 50b.
16 The bridge 50b then examines the packet 165 and determines that it belongs to the hub 26b
17 and passes it on to hub 26b. Hub 26b in turn examines the packet 165 and passes it on to the
18 communicating station 30b.

19 If a communicating station 30 such as the station 30a wants to communicate with a
20 computer or entity 72 outside of the NAN 10, it addresses the packet 165 using the IP address
21 169 of the entity 72. If the outside station 72 wishes to communicate with the station 30a,
22 it also uses an IP address 169 to get into the NAN. This IP address 169 may be either a
23 permanent IP address received from the Internic agency or a masqueraded IP address
24 attributable to the router 62. The outside station 72 sends any return messages using this IP
25 address.

26

1 If the masqueraded IP address is used, the router 62 passes the packet 165 to the
2 switch 22a, which then examines the MAC address 170 without having to refer to the IP
3 address. Thus, one difference between bridges 50 and the routers 62 of the present invention
4 is that a bridge 50 reads only at the MAC level while a router 62 reads at the IP level.

5 The outside station 72 could also be part of a NAN other than the NAN-10. The
6 outside station 72 could communicate using MAC addresses to other outside stations 72
7 within its own NAN, but once it wished to communicate with an entity outside its own NAN
8 such as the communicating station 30a, it then must use an IP address to pass packets 165
9 through the Internet with the use of routers 62.

10 As presently contemplated, each NAN 10 may have 10,000 or more communicating
11 stations 30. A community having more than 10,000 locations wanting to subscribe to the
12 NAN 10 would require more than one NAN 10. Additionally, under the present system, this
13 maximum number may be increased by increasing the speed of the local backbone 16. The
14 speed of the local backbone may be increased up to, for instance, a gigabit per second of
15 throughput without having to reinstall the communicating lines. To increase the number of
16 subscribing communicating stations 30 within a NAN-10, the firmware constituting the
17 software within the client stations server, hubs, bridges and switches are replaced, in an
18 operation that is substantially transparent to the communicating stations 30.

19 Stations within the different NANs preferably communicate with each other over the
20 Internet, as discussed. Nevertheless, within each NAN communications are routerless in the
21 preferred embodiment.

22 Presently, the standard for communications on the inner backbone 16 is 10-base-T,
23 whereas the fiber communications on the fiber backbone 14 are set at 100-base-T. NAN 10
24 communications preferably utilize the Ethernet 802.3 standard which is the standard
25 presently relied upon by most Internet and network organizations. The Ethernet 802.3
26

1 standard is used in one embodiment of the NAN for packet encapsulation for transfer of the
2 packets 165 over communication lines 36, 38.

3 In order for a new communicating station 30 to be admitted to communicate on the
4 NAN 10, it must first establish communications with the server 18. The server 18, as
5 described, maintains a binding between IP addresses and MAC addresses. The client
6 software 126 which is installed on every communicating station 30 provides the
7 communicating station 30 with the proper MAC address of the server 18. Thus the
8 communicating station communicates with the server 18 to receive a localized non-routable
9 IP address for use in communications external to the NAN-10.

10 In one embodiment, the communicating station 30 may be given a permanent IP
11 address issued by Internic or may be given a non-routable address and use the masquerading
12 procedure discussed above. Additionally, there may be several different types of IP
13 addresses issued. As discussed, routable and non-routable IP addresses may be issued as
14 well as filtered IP addresses that filter content received from the Internet. Additionally, an
15 IP address may be partially or fully functional depending on whether the communicating
16 station 30 has paid a monthly or yearly fee.

17 Every station 30 checks in with the server 18 at the initial login in one embodiment,
18 but if the server 18 is not functioning, the stations 30 may still continue to operate with the
19 previously issued IP address. E-mail messages may be sent to a permanent IP address, or
20 may be routed in the manner of outside station 72 communications as discussed above.

21 In addition to the hardware and systems described above, appropriate new hardware,
22 software, and systems may be included in the NAN to enable prioritization of traffic and
23 reduction of broadcast traffic through address caching. Figures 5 through 19 are presented
24 to illustrate such hardware, software, and systems, as well as the methods utilized for traffic
25 prioritization and reduction.

26

1 In the following figures, a number of definitions are relevant. A "data transmission"
2 is simply a digital or analog signal transmitted to a destination. A "packet" 165 is a data
3 transmission bundled in suitable form for delivery over the Internet 34, as depicted in Figure
4 4a. A "switching station" 200 refers to any device that carries and manages data
5 transmissions between multiple ports. Thus, switches 22, hubs 26, repeaters 24, and bridges
6 30 may all be switching stations.

7 However, the packet prioritization and traffic reduction methods of the current
8 invention are well suited to use with switches 22 described above. A "property" of a data
9 transmission is simply any characteristic of the data transmission that can be obtained by a
10 switching station. This includes not just information encoded in the packet 165, but also any
11 other information the switching station could obtain, such as the identification of the port
12 through which the packet 165 entered, characteristics of other packets 165 arriving with the
13 packet 165, etc.

14 "Relative importance" of a data transmission or a packet 165 refers to how important
15 it is that the data transmission reach its destination rapidly. This is relative to the importance
16 of other data transmissions sent through the NAN 10. "Priority" is a related term. The
17 priority of a packet is a designation that determines whether the packet is transmitted before
18 or after other packets. This determination must be made when not all can be simultaneously
19 transmitted, as is the case when multiple packets must go through a single outgoing port.
20 Thus, the relative priorities of multiple packets may be compared to determine a
21 "transmission order" of the packets.

22 Priority may be quantified with a gradation of values, or may be boolean, i.e., "high
23 priority" or "low priority." A "threshold value" may be used to obtain boolean priority with
24 reference to a certain value of the property. For example, if the property of the data
25 transmission is above the threshold value, priority of the data transmission is high, and where
26 the property is equal to or less than the threshold value, priority is low.

1 A "database" is simply an ordered listing of data stored for future retrieval in a
2 memory device. A "destination" of a data transmission or packet 165 refers to an ultimate,
3 terminal destination, rather than to locations of switches en route to the destination. Thus,
4 a destination MAC address is the address to which a packet 165 will ultimately be
5 transmitted. Likewise, an "origin" is a location from which the packet first originated, as
6 opposed to switches upstream of the switching station under analysis. "Location" refers to
7 any origin or destination on the NAN 10, hence, any communication station 30 may be
8 encompassed within the word "location."

9 A "computer-readable medium" is any physical object that can store information in
10 a form directly readable by a computer. Thus, magnetic, optical, and electrical storage
11 devices are all contemplated, as well as any other method of storing information directly
12 accessible to a computer. Hard disks, floppy disks, CD/DVD ROM drives, RAM chips,
13 punch cards, and the like are all examples of computer-readable media. "Instructions" are
14 simply steps to be carried out by a processor, located within a computer-readable medium.
15 The instructions may be provided by hardware, software, firmware, or any suitable
16 combination thereof.

17 A "switching system" includes a switching station and any other components added
18 to improve the quality of switching, such as a prioritization system or traffic reduction
19 system. A prioritization system is an apparatus that acts to assign priorities to data
20 transmissions in a network such as the NAN 10. A traffic reduction system is an apparatus
21 that reduces unnecessary traffic on a network.

22 Referring to Figure 5, one embodiment of a switching system 199 is shown, including
23 a switching station 200 with hardware suitable for packet prioritization and traffic reduction
24 through ARP caching. Packet prioritization and ARP caching may function independently
25 of each other; therefore, a NAN 10 may carry out one method, yet not the other. Different
26 switching stations 200 within a single NAN 10 may be differently configured to carry out

1 packet prioritization, ARP caching, both methods, or neither one. Switching stations 200
2 with few communication stations 30 connected may, for example, derive less benefit from
3 packet prioritization and ARP caching than those with many communication stations 30.

4 The switching station 200 has a plurality of ports 202, preferably from four to twenty-
5 four in number. Each port 202 connects to one of the communication lines 38, which
6 preferably provide full-duplex (i.e., simultaneous, two-way) data transmission. Thus, each
7 port 202 may simultaneously send and receive data. Consequently, the terms "outgoing port"
8 and "incoming port" refer equally to all of the ports 202, and simply delineate what the
9 function of the port 202 is within the process being described. Each port 202 preferably has
10 a buffer 204 to store incoming packets 165 from the port 202. These may be stored in the
11 form of a first-in, first-out (FIFO) stack, so that packets 165 are queued up to be removed
12 from the buffer 204 for processing in the order in which they were received.

13 Each buffer 204 is connected to a bus 206, which operates to transfer data to various
14 components of the switching station 200 at a certain bus speed. A processor 208 connects
15 to the bus 206 to process and manipulate data from the buffers 204. The processor 208 may
16 be of any known type, such as a standard microprocessor, reduced instruction set computing
17 (RISC) processor, field programmable gate array (FPGA), or application-specific integrated
18 circuit (ASIC).

19 A microprocessor is capable of performing a wide variety of instructions, but is not
20 highly specialized to perform any specific instruction set. A RISC processor is more
21 specialized, but is still designed to carry out a comparatively wide variety of instructions.
22 An FPGA is reconfigurable to carry out specific task sets, but is not as fast as an ASIC,
23 which is highly specialized, but not reconfigurable. The ASIC contains a number of logic
24 gates that are fixed in place, and are not reprogrammable. Since the switching station 200
25 will always perform a limited set of instructions, an ASIC is an ideal choice for the processor
26 208. Due to its highly specialized nature, the processor 208 in the form of an ASIC may

1 operate at a speed of 8.4 Gigahertz or greater, thus permitting data transmission through the
2 switching station 200 with very little delay.

3 The processor 208 is connected to a program memory 210, which contains
4 instructions or reference data for the operation of the processor 208. Although many
5 instructions of an ASIC are built into the configuration of gates used, and therefore hard-
6 coded into the processor 208, certain information or instructions needed by the processor 208
7 may be stored in the program memory 210 separate from the processor 208. The program
8 memory 210 may optionally be omitted, if the processor 208 is configured to contain all
9 needed instructions and information.

10 The program memory 210 is situated within a computer-readable medium of any
11 suitable type, such as one or more standard DIMM (Dual In-line Memory Module) or SIMM
12 (Single In-line Memory Module) random access memory (RAM) modules, programmable
13 read-only memory (PROM) modules, electrically erasable PROM (EEPROM) modules,
14 static RAM (SRAM) modules, flash RAM modules, and the like. However, the program
15 memory 210 is preferably of a nonvolatile type, so as to retain information in the event of a
16 loss of electric power to the switching station 200. Additionally, the program memory 210
17 is preferably read-only to avoid any alteration or corruption of information in the program
18 memory 210. Thus, a PROM module or chip is well-suited for use to form the program
19 memory 210.

20 In addition to the program memory 210, the processor 208 may be connected to a
21 multiplexer 211 designed to unify streams of information from simultaneous sources,
22 through interleaving or a similar process. Thus, data from all the buffers 204 and the
23 processor 208 may be transferred into and out of a cache 212 through the multiplexer 211.
24 The multiplexer 211 maybe integrated with the cache 212.

25 The cache 212 is designed to store information pertaining to the operation of the
26 switching station 200. Like the program memory 210, the cache 212 may be embodied as

any suitable memory type such as one or more RAM DIMM or SIMM modules, PROM modules, EEPROM modules, SRAM (Static RAM) modules, flash RAM modules, or the like. Preferably, the cache 212 is erasable, and may be volatile because the information stored in the cache 212 may not be essential to the operation of the switching station 200. SRAM is well adapted for use in the cache 212.

The bus 206 may be connected to an interrupt controller (IC) 219, which permits the switching station 200 to actively interface with a prioritization system 220 connected to work in concert with the switching station 200. The prioritization system 220 preferably takes the form of a packet prioritization station 220. The packet prioritization station 220 may be integrated with the switching station 200, and may even utilize the program memory 210, processor 208, cache 212, multiplexer 211, and bus 206 of the switching station 200. This may be accomplished by providing a new set of instructions in the program memory 210 designed to carry out packet prioritization. However, the packet prioritization station 220 preferably has its own set of independent hardware, so as to be interchangeably usable with any switching station 200, and so as to avoid slowing the operation of the switching station 200. Thus, the packet prioritization station 220 may be located on an auxiliary expansion card or board (AEC), which may be connected to the switching station 200 in modular fashion.

The IC 206 may be integrated with the switching station 200 or the packet prioritization station 220, or may be a separate component from the stations 200, 220. The IC 219 may be connected to a bus 222 located in the packet prioritization station 220 such that data from the bus 206 of the switching station is transmitted to the bus 222 of the packet prioritization station 220. A processor 224, program memory 226, and cache 228 for the packet prioritization station 220 are, in turn, in communication with the bus 222.

As with the processor 208, program memory 210, and cache 212 for the switching station 200, the processor 224, program memory 226, and cache 228 may be of any suitable

1 type. However, the processor 224 is preferably a RISC based processor. This would enable
2 a generalized AEC with a RISC processor to be configured for use as the packet
3 prioritization station. Likewise, the program memory 226 is in a computer-readable medium,
4 preferably comprising an EEPROM module, to enable use of a general-purpose AEC to form
5 the packet prioritization station 220. The EEPROM may then be reconfigured to permit use
6 of the AEC in a different role. The cache 228 is preferably an SRAM module, so as to be
7 erasable and rewritable.

8 The bus 206 is also connected to another interrupt controller (IC) 229, which enables
9 the switching station 200 to interface with a traffic reduction system 230, which may take the
10 form of an ARP caching station 230. As with the packet prioritization station 220, the IC
11 229 may be located in the switching station 200 or the ARP caching station 230. A bus 232
12 in the ARP caching station 230 is in communication with the IC 229 to transmit and receive
13 data from the switching station 200.

14 The ARP caching station, like the packet prioritization station, may be integrated into
15 the switching station 200, and may even operate using the program memory 210, processor
16 208, cache 212, and bus 206 of the switching station 200. However, like the packet
17 prioritization station 220, the ARP caching station 230 is preferably an independent module,
18 which may be located on an AEC. Thus, the processor 234 may be a RISC processor, the
19 program memory 236 may be an EEPROM module, and the cache 238 may be an SRAM
20 module. The ARP caching station 230 and packet prioritization station 220 may thus both
21 operate as independent modules in communication with the switching station 200.

22 Referring to Figure 6, one possible embodiment of one of the buffers 204 of the
23 switching station 200 is shown. Packets 165 are queued in the buffer 204 for FIFO
24 processing. A packet 165 may be substantially as described in connection with Figure 4a.
25 The MAC layer addresses 170 may be located within the header 166 at the periphery of the
26 packet, with separate origin 240 and destination 242 MAC addresses. A designated value

1 in the broadcast field 171 denotes that the packet 165 is to be broadcast throughout all or a
2 specified portion of the NAN 10. An IP address 169 is also provided. Data 167 may be
3 included, or in the case of a packet such as an ARP request, no data need be sent. The footer
4 168 denotes the end of the packet.

5 Referring to Figure 7, one possible embodiment of the program memory 210 of the
6 switching station 200 is depicted. A number of executable modules designed to carry out the
7 method of the current invention may be stored in the program memory 210. As described
8 previously, some or even all of these modules may be hard coded into an ASIC to form the
9 processor 208. However, for purposes of illustration for the following discussion, these
10 instructions are simply represented logically as modules within some form of program
11 memory 210. The modules may be any set of one or more executable instructions to perform
12 a function.

13 A packet reception module 250 handles operations incident to receipt of a packet
14 from one of the ports 202. A cache reading module 252 retrieves information from the cache
15 212 for manipulation by the processor 208. A cache writing module 254, similarly, writes
16 data to the cache 212 for subsequent use. A packet deleting module 256 deletes unnecessary
17 packets from the buffers 204. A port routing module 258 decides which port a given packet
18 should be sent to in order to reach its destination.

19 A comparison module 260 compares separate values or entries to determine whether
20 they are the same, such as comparing an address from the MAC layer 170 of a packet 165
21 with a value in the cache 212 to determine whether the address has been stored in the cache
22 212. A blocking module blocks incoming ports 202 with packets 165 routed to the same
23 outgoing port 202 to permit collisions of packets 165 exiting the switching station 200. The
24 operation of the modules 250, 252, 254, 256, 258, 260, and 262 of the program memory 210
25 will be further clarified by the description of the method of operation of the present
26 invention, to be provided in the description of Figures 13-19.

1 Referring to Figure 8, one possible embodiment of the cache 212 of the switching
2 station 200 is shown. The cache 212 may contain a database 264 in the form of a table 264
3 associating MAC layer 170 addresses with ports 202. A MAC layer 170 address is simply
4 a location identifier, which may act as either an origin MAC address 240, or a destination
5 MAC address 242. Location fields 265 may be provided, in which MAC addresses are
6 stored. The MAC address simply defines a location of a communication station 30 that is
7 accessible from a given port. Port fields 266 corresponding to the location fields 265 show
8 which port 202 a packet 165 must be sent through to reach a given destination MAC address
9 242.

10 Vacant fields 267 in the table 264 may be filled by retrieving the origin MAC address
11 240 from a packet 165 received through a port 202. The origin MAC address 240 is recorded
12 in a vacant location field 268, and the port 202 through which it was received is recorded in
13 a vacant port field 269 corresponding to the field 268. Thus, a MAC layer 170 address is
14 "bound," or associated, with the port 202 through which the MAC layer 170 address is
15 accessible.

16 Packets 165 may then be routed to the appropriate port 202 by the switching station
17 200 by looking up the destination MAC address 242 in the location fields 265 and sending
18 the packet 165 through the corresponding port in the port fields 266. As depicted in Figure
19 8, one port 202 may appear several times in the port fields 266 because there may be
20 switches downstream from the switching station 200, so that one port 202 leads to many
21 unique destination MAC addresses 242.

22 Referring to Figure 9, one possible embodiment of the program memory 226 of the
23 packet prioritization station 220 is shown. As with the switching station 200, instructions
24 may be programmed in any way or even hard-wired into the processor 224 of the packet
25 prioritization station. However, for the sake of discussion, instructions for the processor 224
26 are illustrated logically as modules residing in some form of program memory 226.

1 A cache reading module 270 retrieves data from the cache 228, while a cache writing
2 module 272 stores information in the cache 228 for future retrieval. An incrementing module
3 provides the ability to cyclically analyze the packets 165 in each of the buffers 204 to
4 determine which receives priority for a given outgoing port 202. A marking module 276 is
5 provided to allow the packet prioritization station 220 to mark a given port 202 or packet 165
6 for priority analysis. A comparison module 278 compares separate values or entries to
7 determine whether they are the same, such as comparing an address from the MAC layer 170
8 of a packet 165 with a value in the cache 238 to determine whether the address has been
9 stored in the cache 238. The operation of the modules 270, 272, 274, 276, and 278 will be
10 clarified subsequently, as the method of operation of the present invention is described.

11 Referring to Figure 10, one possible embodiment of the cache 238 of the packet
12 prioritization station 220 is depicted. The cache 238 maintains some property pertaining to
13 incoming packets 165. This property will subsequently be used to prioritize the packets 165
14 for transmission. Preferably, the property includes a number of origin MAC addresses 240
15 that have previously sent transmissions to a destination MAC address 242 of the packet 165.
16 These MAC layer 170 addresses may be stored in the following fashion.

17 A database 290 or table 290 may be stored in the cache 238 to track destination MAC
18 addresses 242 and associate them with origin MAC addresses 240 from which they have
19 received packets 165. Thus, destination fields 292 store destination MAC addresses 242,
20 while several origin fields 294 are associated with each destination field 292. After receiving
21 a packet 165, the packet prioritization station 220 may determine how many origin MAC
22 addresses 240 have previously sent data to the destination MAC address 242 of the packet
23 165 by looking up the destination MAC address 242 in the destination fields 292 and
24 counting the origin MAC addresses 240 in the corresponding origin fields 294.

25 Vacant fields 296 of the table 290 may be filled by storing destination MAC
26 addresses 242 and origin MAC addresses 240 from packets 165 received by the switching

1 station 200. The packet prioritization station 220 may first look up the destination MAC
2 address 242 from an incoming packet 165 in the destination fields 292 to determine whether
3 it is present, and add it if it is not. The packet prioritization station 220 may then add the
4 origin MAC address 240 from the incoming packet 165 to a vacant field corresponding to
5 the destination MAC address 242 among the origin fields 294. Thus, each destination MAC
6 address 240 in the table 290 has at least one origin MAC address 242 associated with it. The
7 number of origin fields 294 may be limited to permit up to a maximum number of origin
8 MAC addresses 240 to be stored for each destination MAC address 242, for example, 16
9 origins per destination.

10 Referring to Figure 11, one possible embodiment of the program memory 236 of the
11 ARP caching station 230 is shown. As with the switching station 200 and the packet
12 prioritization station 220, the executable modules shown in the program memory 236 may
13 be configured in any suitable manner, including being hard-wired into the processor 234.

14 A cache reading module 302 retrieves data from the cache 238. A comparison
15 module 304, like those of the switching station 200 and the packet prioritization station 220,
16 compares two values to determine whether they are the same. For example, the comparison
17 module 304 may compare an address from the MAC layer 170 of a packet 165 with a value
18 in the cache 238 to determine whether the address has been stored in the cache 238. A cache
19 writing module 306 stores data in the cache 238 for subsequent retrieval. A packet
20 preparation module 308 creates a packet in the form of an ARP response of the proper
21 format, to be sent through one of the ports 202 of the switching station 200. The proper
22 format may be whatever packet architecture is currently in use on the NAN 10 for an ARP
23 response. Typically, this is a packet 169 with some special designation to indicate that it is
24 an ARP response. The operation of the modules 302, 304, 306, and 308 will be clarified
25 during the discussion of methods of operation, starting with the description of Figure 13.

26

1 Referring to Figure 12, one possible embodiment of the cache 238 of the ARP
2 caching station 230 is shown. A database 310 in the form of a table 310 may be stored in the
3 cache 238, with IP addresses 169 associated with MAC layer 170 addresses. The table 310
4 may contain bound entries 312 and vacant fields 314 to accept new entries. IP address fields
5 316 store the IP addresses 169, while associated MAC address fields 318 contain MAC layer
6 170 addresses corresponding to the IP addresses 169.

7 When an ARP broadcast is received, it will take the form of a packet 165 with a
8 designated IP address 136 corresponding to a requested MAC layer 170 address sought by
9 the originator of the broadcast. The ARP caching station checks the IP address fields 316
10 to determine whether the designated IP address 136 is stored. If it is found, the ARP caching
11 station may then determine whether the designated IP address 136 has an associated,
12 requested MAC layer 170 address stored in the MAC address fields 316. If so, the ARP
13 caching station returns the requested MAC layer 170 address to the originator of the
14 broadcast. If the MAC layer 170 address is not found in the table 310, the ARP caching
15 station 220 permits propagation of the packet 165 containing the ARP broadcast through the
16 ports 202 of the switching station 200.

17 The vacant fields 314 may be filled by storing IP addresses 169 and MAC layer 170
18 addresses from incoming packets 165. For example, when a packet 165 is received by the
19 switching station 200, the ARP caching station 230 may read the IP address 136 from the
20 packet 165 and look it up in the IP address fields 316 to see if it has been stored. If the IP
21 address 136 has been stored, the ARP caching station 220 checks the MAC address fields
22 318 to determine whether a corresponding MAC layer 170 address has been recorded in the
23 table 310. If necessary, the ARP caching station 230 adds the MAC layer 170 address to the
24 associated field in the MAC address fields 318. If the IP address 136 has not been stored,
25 the ARP caching station 220 stores it in a vacant IP address field 320. The ARP caching
26

1 station then adds the MAC layer 170 address to the corresponding vacant MAC address field
2 322.

3 Referring to Figure 13, one embodiment of an overall method 330 of handling a
4 packet according to the invention is shown. Steps and queries may be added, deleted, or
5 rearranged, as suited to the characteristics of the NAN 10. Several of the steps of the
6 following description will be described in greater detail in Figures 14-19.

7 In a preliminary processing step 332, a packet 165 is received and preliminarily
8 processed by the switching station 200. Then, either the processor 208 of the switching
9 station 200 or the interrupt controller 219 executes an availability test 334. The availability
10 test 334 determines whether the packet prioritization station 220 is available, or not. This
11 step of the method 330 is necessary because a highly-specialized ASIC-based processor 208
12 operates at comparatively high speed, on the order of 8.4 Gigahertz. A RISC-based
13 processor 224 of a packet prioritization station 220, on the other hand, may function at
14 around 200 Megahertz, a speed orders of magnitude lower than that of the ASIC-based
15 processor 208.

16 Thus, the packet prioritization station 220 may only be available to accept data during
17 certain cycles of the switching station 200. If the packet prioritization station 220 is
18 available, packet information, such as MAC layer 170 addresses, will be transmitted to the
19 packet prioritization station 220 from the switching station 200. The interrupt controller 219
20 preferably produces an intransitive interrupt, i.e., an interrupt that continues with the primary
21 process regardless of the operation of the auxiliary process, for the availability test 334.
22 Thus, the switching station 200 continues processing of the packet 165 whether or not data
23 has been sent to the packet prioritization station 220.

24 When available, the packet prioritization station 220 receives and stores the origin
25 and destination MAC addresses 240, 242 from the packet 165 in a priority information
26 storage step 336. The packet prioritization station 220 need not store MAC layer 170

1 addresses from every single packet 165 received; a representative sampling is sufficient to
2 properly prioritize outgoing packets 165 later in the process 330. An intransitive interrupt
3 permits the packet prioritization station to obtain such a representative sampling without
4 slowing operation of the switching station 200.

5 After the availability test 334, a broadcast test 338 may be performed by the processor
6 208, but is preferably carried out by the interrupt controller 229 in communication with the
7 ARP caching station 230. The broadcast test 338 determines whether the packet 165 is a
8 broadcast. The broadcast test 338 is preferably of a transitive type, since the status of the
9 packet 165 must be resolved before operation of the switching station 200 may continue.
10 Thus, the broadcast test 338 interrupts the operation of the switching station 200, if
11 necessary, to process broadcast packets 165.

12 A broadcast packet 165 may have a specially designated destination MAC address
13 242, an empty destination MAC address 242, or a specially designated broadcast field 171.
14 If the packet 165 is a broadcast, an ARP request test 340 is executed by the interrupt
15 controller 229, or preferably by the processor 236 of the ARP caching station 230. If
16 executed by the processor 236, no interrupt occurs because the switching station 200 is still
17 waiting for the status of the packet 165 to be determined. The ARP request test 340
18 determines whether the packet 165 is an ARP broadcast, or a broadcast requesting a
19 requested MAC layer 170 address for a designated IP address 136. Special designations in
20 the MAC layer 170 addresses, IP address 136, or data 167 of the packet 165 may be read to
21 make this determination.

22 If the packet 165 is an ARP broadcast, the request is then processed by the ARP
23 caching station 230 in an ARP request processing step 341. This entails creating a response
24 with the requested MAC layer 170 address if the requested MAC layer 170 address is in the
25 cache 238 of the ARP caching station 230. Otherwise, the ARP caching station 230 permits
26 the ARP broadcast packet 165 to be broadcast.

1 If the packet 165 is a broadcast, but is not an ARP broadcast, it need not have further
2 interaction with the ARP caching station, because it has no associated IP address 136 and
3 MAC layer 170 address to store, and does not require an ARP response. Thus, the packet
4 165 proceeds to a packet routing step 342 in which the packet 165 returns to the switching
5 station 200 for routing.

6 "Routing" generally refers to the process of selecting a path for a data transmission.
7 In the context of Figures 5 through 19, "routing" refers more specifically to determining
8 which of the ports 202 a packet 165 should be sent through to reach a given destination.
9 "Allocation" is simply the process of assigning a packet 165 a destination MAC address 242
10 or an IP address 136 denoting a destination. This may be done by using a destination MAC
11 address 242 contained within the packet 165. However, significant benefits may be obtained
12 through the use of additional steps to determine where the packet 165 should most efficiently
13 be sent, as described in greater detail below.

14 If the broadcast test 338 determines that the packet 165 is not a broadcast, the packet
15 165 will be processed by the ARP caching station 230 in an ARP caching step 344. In the
16 ARP caching step 344, the IP address 136 and destination MAC address 242 are stored in
17 the cache 238 of the ARP caching station in associated form for future use. The packet 165
18 is then routed by the switching station 200 in the packet routing step 342.

19 After routing, yet a multiple routed packets test 345 is executed, possibly by the
20 processor 208 of the switching station 200, but preferably by the interrupt controller 219
21 linked to the packet prioritization station 220. The multiple routed packets test 345
22 determines whether multiple packets 165 in the buffers 204 have been routed to a single
23 outgoing port 202. Like the broadcast test 338, the multiple routed packets test 345
24 preferably takes the form of a transitive interrupt, because the switching station 200 cannot
25 proceed to block ports 202 until an order for blocking has been determined. If multiple
26 packets 165 are routed to a single outgoing port 202, a blocking decision step 346 occurs in

1 which a blocking decision is made by the packet prioritization station 220 to determine
2 which packet 165 is sent first.

3 The blocking decision of the blocking decision step 346 may be made by assigning
4 a high priority to packets 165 being sent to destination MAC addresses 242 with more than
5 a threshold number of associated origin MAC addresses 240 in the cache 228 of the packet
6 prioritization station 220. The remaining packets receive a low priority. Although multiple
7 priority gradations may be used, high and low are simple and enable rapid operation of the
8 packet prioritization station 220. Among packets 165 with the same priority, the blocking
9 decision step 346 may unblock ports 202 in a round robin, or cyclical form. Unblocking is
10 simply the process of permitting the first queued packet 165 in a buffer 204 to exit through
11 its outgoing port 202, or ports 202, in the case of a broadcast.

12 If multiple packets 165 are not routed to a single outgoing port 202, no special
13 blocking decision need be made. Thus, in a round robin blocking step 348, no priority need
14 be assigned to any packet 165, but unblocking of the ports 202 occurs in round robin, or
15 cyclical form among all packets. Finally, after one or more ports 202 has been unblocked,
16 the packet 165 is transmitted through the port 202 or ports 202 in a packet sending step 350.
17 The process 330 then begins anew with the next packet 165.

18 Referring to Figure 14, the preliminary processing step 332 is shown in greater detail.
19 In a packet receiving step 360, a packet 165 is received through an incoming port 202 and
20 enters the associated buffer 204. In an origin cached test 362, the switching station 200
21 determines whether the origin MAC address 240 is in the location fields 265 of the cache
22 212. If not, the switching station 200 performs a port association step 364.

23 The port association step 364 may include storage of the origin MAC address 240 in
24 the vacant location field 268, and storage of an identifier (such as a letter) for the incoming
25 port 202 in the vacant port field 269. The switching station 200 will then be able to route
26

1 response packets 165 back to that origin MAC address 240 without broadcasting the packets
2 165 through multiple ports 202.

3 Referring to Figure 15, the priority information storage step 336 is shown in greater
4 detail. In a destination cached test 370, the packet prioritization station 220 determines
5 whether the destination MAC address 242 of the packet 165 is in the cache 228. If the
6 destination MAC address 242 is not found in the cache 228, the destination MAC address
7 242 is added to the destination field 292 of the cache 228 in a destination storage step 372.
8 There is no need to proceed further, so the packet prioritization station 220 again becomes
9 available in an availability step 373.

10 If the destination MAC address 242 was found in the cache 228, an origin associated
11 test 374 determines whether the origin MAC address 240 of the packet 165 has been
12 associated with the destination MAC address 242 in the origin fields 294 of the cache 228.
13 If not, in a vacancy test 376, the packet prioritization station 220 checks to see if there is
14 vacancy in the origin fields 294 associated with the destination MAC address 242.

15 If there is vacancy, the packet prioritization station 220 performs an origin storage
16 step 378. In the origin storage step 378, the origin MAC address 240 is added to the
17 appropriate field of the origin fields 294 for the destination MAC address 242. The packet
18 prioritization station 220 then becomes available again in the availability step 373. This also
19 occurs if the origin MAC address 240 is already in the cache 228, or if there is no vacancy.

20 Referring to Figure 16, the ARP request processing step 341 is shown in greater
21 detail. In an IP address cached test 390, the ARP caching station 230 determines whether the
22 designated IP address 136 of the packet 165 is in the IP address fields 316 of the cache 238.
23 If the designated IP address 136 is not found, the designated IP address 136 may be cached
24 in a vacant IP address field 320 of the cache 238 in an IP address caching step 392. If the
25 designated IP address 136 is already present, the ARP caching station 230 performs an IP

26

1 address bound test 394 to determine whether a MAC layer 170 address is bound to the IP
2 address 136.

3 If no associated MAC layer 170 address is found, or if the IP address 136 was just
4 cached in the IP address caching step 392, the ARP caching station 230 performs an ARP
5 request allocating step 396. In the ARP request allocating step 396, the ARP broadcast 165
6 is allocated to all ports 202 of the switching station except the incoming port 202. In effect,
7 since the cache 238 does not contain the requested MAC layer 170 address, the ARP
8 broadcast 165 is allocated for further broadcast from the switching station 200.

9 If the requested MAC layer 170 address is found in the cache 238, the ARP caching
10 station 230 initiates a response creation step 398, in which a response to the ARP broadcast
11 165 is created. The response may take the form of a packet 165 with the origin MAC address
12 240 of the ARP broadcast used to form the destination MAC address 242 of the packet 165
13 of the response. The packet 165 of the response has thereby been allocated to a single
14 destination, and will only have to be sent through a single port 202. The requested MAC
15 layer 170 address is contained in the packet 165 of the response, either as the origin MAC
16 address 240, or in the data 167 of the packet 165. Thus, the response containing the
17 requested MAC layer 170 address is returned directly to the originator of the ARP broadcast.

18 Referring to Figure 17, the ARP caching step 344 is shown in greater detail. Since
19 the packet 165 is not a broadcast, as determined by the broadcast test 338, it must have a
20 destination MAC address 242. Consequently, the ARP caching station 230 may perform a
21 MAC address binding step 402. The MAC address binding step 402 entails adding the
22 destination MAC address 242 to the appropriate field of the MAC address fields 318 to bind
23 it to the IP address 169 of the packet 165.

24 The IP address 169 was previously cached in the IP address caching step 392. Thus,
25 the destination MAC address 242 may be obtained from the cache 238 of the ARP caching
26 station 230 for response to another, subsequently received packet 165 containing an ARP

1 request. In a MAC response allocation step 404, the packet 165 may simply be assigned to
2 the destination MAC address 242 from the packet 165.

3 The cache 238 is preferably cleared periodically. Since IP addresses 169 from most
4 ISP's are only temporary or semi-permanent, a user logging onto an Internet service provider
5 (ISP) for a second time may have a different IP address than that of the prior session. Thus,
6 clearing the cache 238 may prevent inaccuracies from building up and slowing down the
7 NAN 10. Clearing the cache 238 periodically also permits a smaller cache 238 to be used.
8 Clearing may take place after a suitable time period, such as one day.

9 Referring to Figure 18, the packet routing step 342 is shown in greater detail.
10 Routing may consist of adding a tag or identifier (not shown) to the packet 165 in the buffer
11 204, storing a suitable port-to-packet correlation (not shown) in the cache 212, or any other
12 method of linking one or more ports 202 to the packet 165. In a destination cached test 410,
13 the switching station 200 determines whether the destination MAC address 242 is in the
14 cache 212 of the switching station 200.

15 This may be accomplished by looking up the destination MAC address 242 in the
16 MAC address fields 265 of the cache 212. If the destination MAC address 242 is not found,
17 the switching station 200 has no record of which port 202 leads to the destination MAC
18 address 242, and must therefore route the packet 165 to all ports 202 except the incoming
19 port 202 in an all ports routing step 412.

20 If the destination MAC address 242 is found in the cache 212, the switching station
21 200 may then determine whether the destination MAC address 242 is associated with the
22 incoming port 202 in a destination associated test 414. Thus, the switching station 200 may
23 be configured to check the field of the port fields 266 that corresponds with the destination
24 MAC address 242 from the MAC address fields 265.

25 If the port 202 associated with the destination MAC address 242 is the incoming port
26 202 of the packet 165, the packet 165 is already travelling through the lines and switches

downstream of the port 202 through which it needs to be sent, so the packet 165 need not be sent at all. Thus, if the destination MAC address 242 is associated with the incoming port 202, the packet 165 is deleted from its buffer 204 in a packet deleting step 416. If the destination MAC address 242 is associated with a different port 202 than the incoming port 202, switching station 200 performs an associated port routing step 418, in which the packet 165 is routed to the associated port 202.

Referring to Figure 19, the blocking decision step 346 is shown in greater detail. Since the packet routing step 342 described previously occurs for each buffer 204, several packets 165 have been routed to their appropriate ports 202. If packets 165 from two or more buffers 204 are routed to a single port 202 simultaneously, the switching station 200 will need to block all but one of the buffers 204 to transmit a single packet 165 at a time. This must occur in sequence, until multiple buffers 204 no longer contain packets 165 routed to the same outgoing port 202.

The process followed by the blocking decision step 346 ensures that the blocking decision is made intelligently. When a blocking decision must be made, more important packets 165 are prioritized for transmission. Such a decision may occur according to the process shown in Figure 19. First, in a step 419, the starting port 202 is incremented and marked. The starting port 202 is the port 202 connected to the buffer 204 through which the last transmission was sent. “Incrementing” entails choosing the next port.

Incrementing may be carried out in an arbitrary, cyclical order, for example, W, then X, then Y, then Z, then W again, and so on, for the ports 202 shown in Figure 5. If a data transmission was just sent from the buffer 204 attached to a port X, incrementing the current port 202 causes a port Y to become the current port 202. Simply incrementing the ports 202 in such a cyclical fashion, with no variation to account for priority, may be referred to as “cyclical,” “round robin,” or “alternating” transmission of packets 165.

1 If no port 202 has priority, the packet 165 in the current port 202 is sent. Priority
2 analysis may first be undertaken to determine whether another of the ports 202 should have
3 priority over the current port 202. Thus, in an incrementing and marking step 420, a current
4 port 420 is designated and set to be the same port 202 as the starting port 202. The current
5 port 202 is the port 202 under prioritization analysis.

6 Analysis begins in a port associated test 430, in which the packet prioritization station
7 220 determines whether the destination MAC address 242 of the current port 202 has been
8 routed to only a single outgoing port 202. Thus, in the packet routing step 342, if the packet
9 165 was routed to all ports 202 except the incoming port 202, as in the all ports routing step
10 412, the answer to the port associated test 430 will be "no." If, in the packet routing step
11 342, the packet 165 was routed to a single port 202, as in the associated port routing step
12 418, the port associated test 430 will return a "yes."

13 If the answer is "no," i.e., the packet 165 in the buffer 204 of the current port 202 is
14 routed to multiple ports 202, the current port 202 is incremented to the next port 202. The
15 net effect of the port associated test 430 is to pass over packets 165 that must be broadcast
16 to multiple ports 202 to prioritize packets 165 with a known outgoing port 202. As described
17 above, any type of broadcast uses a comparatively greater portion of bandwidth because it
18 must be sent along multiple routes. Thus, broadcast traffic is delayed by the step 430 in
19 favor of traffic that requires less bandwidth for transmission.

20 After the current port 202 has been incremented, i.e., set to the next port 202 in the
21 cycle, a cycle completed test 440 inquires whether the current port 202 has become the same
22 as the starting port 202. If so, the buffer 204 of the current port 202 is unblocked for
23 transmission through its routed outgoing port 202 or ports 202 in a starting port unblocking
24 step 442. If not, the new current port 202 is analyzed by the step 430. The effect of the cycle
25 completed test 440 is to allow priority analysis to occur for each port 202 only once before
26 a transmission is made. If no port 202 meets the qualifications for priority, the starting port

1 202, i.e., the next port 202 in line after the previous transmission is sent, may be unblocked
2 by the starting port unblocking step 442.

3 If the answer to the port associated test 430 was "yes," i.e., the packet 165 in the
4 buffer 204 of the current port 202 is routed to a single outgoing port 202, priority analysis
5 continues on the current port 202 in a step 450. In the step 450, the packet prioritization
6 station 220 determines whether four or more origin MAC addresses 240 are associated, or
7 bound, to the destination MAC address 242 of the packet 165 in the buffer 204 of the current
8 port 202. This is accomplished by looking up the destination MAC address 242 of the packet
9 165 in the destination fields 292 of the cache 228, and counting the origin MAC addresses
10 240 in the origin fields 294 associated with the destination MAC address 242. If more than
11 some threshold number, for example, four, origin MAC addresses 240 are associated with
12 the destination MAC address 242, the current port 202 may be unblocked to send the packet
13 165 in a current port unblocking step 452. Otherwise, the packet 165 does not receive
14 priority and the current port 202 is incremented in the current port incrementing step 432 to
15 continue with priority analysis.

16 The effect of the multiple bound origins test 450 is to prioritize packets 165 to
17 destination MAC addresses 242 that have received packets 165 from multiple origin MAC
18 addresses 240. This is effective because communication stations 30 that receive traffic from
19 many locations have been shown to be more likely to be receiving more time-critical traffic,
20 or to have many users. Communication stations 30 that receive data from only a few sources
21 have been shown to be more likely transferring larger amounts of data, for which some delay
22 is acceptable. Thus, the multiple bound origins test 450, with the aid of the cache 228
23 maintained by the packet prioritization station 220, effectively prioritizes transmission of the
24 most important information.

25 "Unblocking" a port 202 enables transmission of the packet 165 in the buffer 204 of
26 that port 202, through its routed outgoing port 202 or ports 202. The cycle described above

1 occurs as many times as necessary to clear the traffic routed to one outgoing port 202. When
2 this has been accomplished, unblocking may simply occur in round robin form, i.e.,
3 cyclically unblocking ports 202 with no priority decision, as in the round robin blocking step
4 348, until the need once again arises to make a blocking decision.

5 As with the cache 238, the cache 228 is preferably cleared periodically. This may be
6 necessary primarily because the usage patterns of a communication station 30 located at a
7 given MAC layer 170 address may change over time. A communication station 30 may be
8 used for a highly time-critical application one day, and then for less critical applications the
9 next day. Clearing the cache 228 effectively resets the priority of communication stations
10 30 on the NAN 10 so that a newer and more accurate determination can be periodically
11 made. Clearing the cache 228 also enables a smaller cache 228 to be used, because fewer
12 MAC layer 170 addresses need be stored. The cache 228 may be cleared after a period of
13 suitable length, such as one day.

14 The NAN of the present invention provides certain advantages including providing
15 high speed (high band width) Internet access at a low cost compared to conventional
16 technologies. Advantages of the NAN also include the capability of real-time video
17 conferencing. The NAN allows a region such as a geographical region of otherwise
18 unrelated entities, such as a town or neighborhood, to be networked in high speed computer
19 communication.

20 The NAN may be financed at least partially by utilities in order to expedite
21 installation and may rely on the rights of way of public utilities such as power companies.
22 The "last mile" dilemma is also solved under the present invention, as the system allows for
23 inexpensive installation of facilities for the "last mile" of a network infrastructure and
24 relatively faster operation thereof. Thus, an advantage of the NAN is that it provides cost
25 effective last mile service and delivery.

26

1 The NAN also operates at very high speeds. Preferably, message traffic is directly
2 delivered to its destination, rather than passing the message traffic through a central server
3 or router. Indeed, in certain embodiments, the NAN efficiencies are achieved without a
4 central server altogether.

5 Additionally, the NAN provides support for a broader variety of devices and types
6 of devices to be networked. The NAN system of the present invention does not rely on the
7 telephone line infrastructure, and consequently eliminates handling errors that occur with
8 user log ons. Additionally, the telephone lines and other telecommunications infrastructure
9 receive less traffic and are less likely to be jammed with message traffic when the NAN is
10 employed to relieve them of being overburdened. Indeed, the NAN in one embodiment
11 achieves total independence from the telecommunication infrastructure.

12 Also, no modem hardware or protocol is necessary at the user facility. Conventional
13 T-1 lines, fiber converters, and cable modems are unnecessary in achieving the much higher
14 speeds of the NAN of the present invention. Additionally, Internet access may be provided
15 over the NAN, and Internet connection may operate at comparatively high speeds. For
16 instance, Internet access may in one example be as high as ten Mbps while employing certain
17 currently available hardware.

18 The NAN allows free competition among Internet service providers and allows them
19 to freely hook into the NAN system. The Internet connectivity is always on and continuous
20 at any given communicating station without the need of a dial-up. Due to the elimination of
21 modems in connecting to the Internet, low data losses are experienced. For instance, hand
22 shaking errors between modems and error data that otherwise arises between modems may
23 be reduced or eliminated. This is largely due to the absence of protocol conversions with the
24 inventive system.

25 The operational hardware and software of the NAN include hubs, packets, bridges,
26 and gateways disposed at different points to allow directly routed, packeted traffic. The

1 system completes routing and distributes traffic at the lowest possible segment. Direct
2 routing may be peer-to-peer rather than being controlled by a switchboard, server, or central
3 office. The results of this arrangement is very high speed packet transfer.

4 The system may rely on MAC addresses and static, masqueraded, IP addressing rather
5 than dynamic IP addressing. The system may provide a binding between a hardware device
6 and a user so the system stores the user's public IP addresses.

7 Additionally, communications within the network are secure and the network is user
8 friendly. The high-speed networking supports real-time communications with cameras.
9 Indeed, because of the low cost, users can connect to more devices, one example of which
10 is utility meters. The system makes remote meter reading and monitoring of other types of
11 utility services cost effective.

12 The NAN of the present invention is also unique in that no network administration
13 is necessary to control local message traffic. Traffic may be independent of any governing
14 authority. Additionally, because the Internet is both a large scale system and localized within
15 a geographic area, business services such as advertising can be offered locally, making them
16 more efficient. Thus, local advertising may be directed to a local audience. The system may
17 support interconnection with virtually any devices within a community. The system may
18 utilize permanent IP addresses due to a unique Dynamic Host Configuration Protocol
19 (DHCP).

20 The NAN 10 of the present invention is further distinguished from the prior art in that
21 packet prioritization is provided for packets transmitted through the switching stations 200
22 of the NAN 10. The switching stations 200 may prioritize traffic to destinations receiving
23 traffic from multiple origins. This accords a generally higher priority to traffic with a higher
24 likelihood of being time-critical, such that packets with a higher relative importance are
25 transmitted first. All of this may be accomplished through the use of packet prioritization
26 stations 220 that can be added or modified at will for use with the switching stations 200.

1. The NAN 10 is further unique in that a method for reducing traffic from ARP
2 broadcasts is provided. This may be accomplished by caching MAC layer 170 addresses and
3 associating them with their corresponding IP addresses. ARP broadcast traffic is reduced by
4 simply returning the requested MAC layer 170 address from the cache 238. This saves a
5 great deal of bandwidth over broadcasting multiple ARP request packets 165 while waiting
6 for a response from the communication station 30 that has the requested MAC layer 170
7 address.

8 The present invention may be embodied in other specific forms without departing
9 from its essential characteristics. The described embodiments are to be considered in all
10 respects only as illustrative and not restrictive. The scope of the invention is, therefore,
11 indicated by the appended claims rather than by the foregoing description. All changes
12 which come within the meaning and range of equivalency of the claims are to be embraced
13 within their scope.

14 What is claimed and desired to be secured by United States Letters Patent is:
15
16
17
18
19
20
21
22
23
24
25
26